

Exact algorithm for the bottleneck 2-connected k -Steiner network problem*

Marcus Brazil[†], Charl J. Ras[‡], Doreen A. Thomas[‡]

Abstract

We present the first exact algorithm for constructing minimum bottleneck 2-connected Steiner networks containing at most k Steiner points, where $k > 2$ is a constant integer. The objective of the problem is – given a set of n terminals embedded in the Euclidean plane – to find the locations of the Steiner points, and the topology of a 2-connected graph N_k spanning the Steiner points and the terminals, such that the length of the bottleneck (the longest edge of N_k) is minimised. The problem is motivated by the modelling of relay-augmentation for optimisation of energy consumption in wireless transmission networks. Our algorithm employs Voronoi diagrams and properties of block cut-vertex decompositions of graphs to find an optimal solution in $O(h(k) n^k \log^{5k-1} n)$ steps, where $h(k)$ is a function of k only.

Keywords: bottleneck optimisation, Steiner network, 2-connected, block cut-vertex decomposition, exact algorithm, wireless networks

1 Introduction

In the design of survivable communication networks, especially wireless sensor networks, the *bottleneck* edge (or longest link) is often considered a critical parameter. This is due to the fact that, in general, maximum transmission power is utilised at the nodes communicating across the bottleneck. Relay-augmentation – the process of introducing (non-sensing) relays in strategic locations relative to the nodes of the original deployment – is one of the most effective methods of reducing the bottleneck parameter, and can be modelled by means of *bottleneck Steiner networks*, where the Steiner points correspond to the relays.

Exact algorithms for constructing optimal survivable (i.e., multiply-connected) bottleneck Steiner networks are absent from the literature, despite the fact that such algorithms (and the machinery and insights upon which their construction depends) are essential for the ultimate goal of effectively modelling multi-parameter survivable network relay-augmentation

*This research was supported by an ARC Discovery Grant

[†]Department of Electrical and Electronic Engineering, The University of Melbourne, Victoria 3010, Australia. email: brazil@unimelb.edu.au; cjas@unimelb.edu.au

[‡]Department of Mechanical Engineering, The University of Melbourne, Victoria 3010, Australia. email: doreen.thomas@unimelb.edu.au

problems. We therefore introduce a new exact algorithm for solving the *bottleneck 2-connected k -Steiner network problem* on a set X of n given nodes embedded in the Euclidean plane, where an optimal solution (minimising the longest edge) is required to be a 2-connected spanning graph containing at most k Steiner points. An upper bound on the number of Steiner points is not only realistic (due to the cost of introducing individual relays) but also necessary to ensure that a solution exists (see [5]).

A naive algorithm for solving the above problem would require more than $O(n^{3k})$ steps, since we may assume that each Steiner point s is at the centre of the circumcircle on exactly two or three neighbours of s . But suppose we are given the subgraph G of an optimal solution N_k , where G is induced by X . Since the degree of any Steiner point can be assumed to be at most 5 in N_k , there are at most $h_0(k)$ topologies \mathcal{N} induced by the Steiner edges of N_k , where $h_0(k)$ is some function of k only. Furthermore, we show that \mathcal{N} may be assumed to be a forest and that it is possible to greatly reduce the number of subsets of X that potentially contain the leaves of \mathcal{N} . We prove the latter by constructing a set $\{J(e_i)\}$ of at most $5k$ labelled subsets of X , where $\{e_i\}$ is the set of leaf-edges of \mathcal{N} , such that if $x_i \in J(e_i)$ is the non-Steiner end-vertex of e_i then, regardless of the choice of x_i , the graph $G \cup \mathcal{N}$ is 2-connected. Since, as we show by using connectivity properties of block cut-vertex decompositions of graphs, an optimal $\{J(e_i)\}$ can be found in a time of $O(\ln^{5k-2} n)$, and an optimal x_i for each $J(e_i)$ along with the optimal locations for all Steiner points can be found in a time of $O(n^k)$ (using a modified version of a method introduced by Bae et al. in [2]), we attain a complexity of $O(n^k \log^{5k-2} n)$ for any given G . A simple binary search, requiring $O(\log n)$ time, finds the correct G , giving a total complexity of $O(n^k \log^{5k-1} n)$. Note that the above-mentioned complexity bounds ignore factors that are functions of k only, and we maintain this convention throughout the rest of the paper.

Our results may be seen as a continuation and generalisation of the algorithms by Brazil et al. in [6], where the same problem was solved for $k = 1, 2$. Besides [6] and a few papers on constructing so called *bottleneck biconnected spanning subgraphs* (see [7, 10, 11, 12, 15]), few results in the literature are directly relevant to our problem. However, much progress has been made in the field of *bottleneck Steiner trees* (that is, when the solution is simply required to be connected). This problem has been shown to be NP-hard in the Euclidean and rectilinear planes, and also in general graphs (see [4, 13, 16]); and recently Bae et al. in [2, 3] and Brazil et al. in [5] have produced exact algorithms that run in polynomial time for constant k .

Section 2 introduces notation and preliminary results, most of which can be found in more detail in [6]. In Section 3 we present the *fixed topology on subsets problem*, and show that the methods of Bae et al. effectively solve it. Fundamental connectivity properties of 2-connected graphs, which we introduce in Section 4, lead to an efficient method of listing all potentially optimal solution topologies. Section 5 then presents our main algorithm.

2 Preliminaries

Throughout this paper we only consider finite, simple, and undirected graphs. A graph G is *connected* if there exists a path connecting any pair of vertices in G . An *isolated component* is a maximal (by inclusion) connected subgraph. A *cut-set* A of G is any set of vertices such that $G - A$ has strictly more isolated components than G ; if $|A| = 1$ then A is a *cut-vertex*. Set A *separates* W from Z in G , where W, Z are subgraphs of G , if every path connecting a vertex of W to a vertex of Z contains a vertex of A . If A separates any subgraphs of G then A is a cut-set of G .

The *vertex-connectivity* or simply *connectivity* $c = c(G)$ of a graph G is the minimum number of vertices whose removal results in a disconnected or trivial graph. Therefore c is the minimum cardinality of a cut-set of G if G is connected but not complete; $c = 0$ if G is disconnected; and $c = n - 1$ if $G = K_n$, where K_n is the complete graph on n vertices. A graph G is said to be c' -*connected* if $c(G) \geq c'$ for some non-negative integer c' . In this paper we make an exception for the connectivity definitions of K_1, K_2 : we assume that $c(K_1) = c(K_2) = 2$.

A *critical edge* of a 2-connected graph is an edge such that its removal reduces the graph's connectivity. From [8] we know that an edge is critical if and only if it is not a chord of any cycle. A *block* is a maximal 2-connected subgraph. For any graph G we denote the longest edge of G (where ties have been broken) by $e_{\max}(G)$ and its length by $\ell_{\max}(G)$. Let X be a set of vertices (called *terminals*) embedded in \mathbb{R}^2 .

Definition 2.1 *The Euclidean bottleneck c -connected k -Steiner network problem requires one to construct a c -connected network N_k spanning X and a set $S_{k'}$ of $k' \leq k$ Steiner points, such that $\ell_{\max}(N_k)$ is a minimum across all such networks. The variables are k' , the set $S_{k'}$, and the topology of the network.*

An optimal solution to the problem is called a *minimum bottleneck c -connected k -Steiner network*, or (c, k) -MBSN. In this paper we focus on the case $c = 2$ with $k \geq 3$. We also assume throughout that $|X| = n \geq 2$.

Let $\{E_i\}$ be a partition of $E(G)$ such that each E_i induces a block Y_i of G , and let $\mathcal{Y}(G) = \{Y_i\}$. Note that each non-cut-vertex of G is contained in exactly one of the Y_i ; each cut-vertex of G occurs at least twice amongst the Y_i ; and for each $i, j, i \neq j$, $V(Y_i) \cap V(Y_j)$ consists of at most one vertex, and this vertex (if it exists) is a cut vertex of G . If Y_i contains exactly one cut-vertex of G then Y_i is a *leaf block*. An *isolated block* contains no cut-vertices of G , i.e., it is a 2-connected isolated component of G . We use $\mathcal{Y}_0(G)$ to denote the set of leaf blocks of G . The *interior* of block Y_i , denoted $\text{int}(Y_i)$, is the set of all vertices of Y_i that are not cut-vertices of G . The unique cut-vertex of G belonging to $Y_i \in \mathcal{Y}_0(G)$ is denoted by $\tau(Y_i)$.

The *block cut forest* (BCF) of G is a forest $F_{\mathcal{Y}(G)}$ with $V(F_{\mathcal{Y}(G)}) = \{Y_i \in \mathcal{Y}(G)\} \cup \{z_i : z_i \text{ is a cut-vertex of } G\}$ and $E(F_{\mathcal{Y}(G)}) = \{Y_i z_j : z_j \in Y_i\}$. For any subgraph F of $F_{\mathcal{Y}(G)}$ the corresponding (vertex-induced) subgraph of G is the graph \hat{F} where $u \in V(\hat{F})$ if and

only if $Y_i \in V(F)$ and $u \in Y_i$; in particular, $G = \widehat{F_{\mathcal{Y}(G)}}$. In Fig. 1 we show three graphs: the first is a graph M of connectivity 1; the second is the graph $F_{\mathcal{Y}(M)}$, where large double circles represent the vertices Y_i and small single circles represent the z_i ; and the third is a simplified depiction of $F_{\mathcal{Y}(M)}$, where degree-two z_i are not shown and edges are drawn as double lines, which we utilise in this paper to avoid ambiguity whenever illustrating graphs containing edges of G and $F_{\mathcal{Y}(G)}$.




Figure 1: The graphs M , $F_{\mathcal{Y}(M)}$, and the simplified depiction of $F_{\mathcal{Y}(M)}$

Theorem 2.2 (see [14]) *The BCF of a graph G with m edges can be constructed in time $O(m)$. As part of the construction we can calculate the connectivity of G , and also specify all cut-vertices and the blocks that contain each cut-vertex.*

Proposition 2.3 ([8]) *If e is a critical edge of any 2-connected graph A then the BCF of $A - e$ is a path containing at least two distinct vertices.*

The following converse to this proposition – used implicitly in many proofs of this paper – is easy to show.

Lemma 2.4 *Suppose that the BCF of a graph A is a path Y_1^A, \dots, Y_q^A with $q \geq 2$, and let $e = xy$ be any edge not in $E(A)$ such that $x \in \text{int}(Y_1^A)$ and $y \in \text{int}(Y_q^A)$. Then $A + e$ is 2-connected.*

We define a counter, $b(\cdot)$, as follows. Let $\{G_i\}$ be the set of isolated components of G . If G_i is an isolated block then let $b(G_i) = 2$, else let $b(G_i) = |\mathcal{Y}_0(G_i)|$. Finally, let $b(G) = \sum b(G_i)$. Essentially $b(G)$ is the number of leaf blocks plus twice the number of isolated blocks occurring in G (recall that isolated vertices and isolated edges are blocks according to our definition).

Lemma 2.5 ([6]) *If G_1 is an edge subgraph of G_2 then $b(G_1) \geq b(G_2)$.*

Let e be any edge of a plane embedded graph. The *lune specified by e* is the region of intersection of the two circles of radius $|e|$ centred at the endpoints of e .

Definition 2.6 (see [7]) *The 2-relative neighbourhood graph on X (or 2-RNG) is the graph R such that $e \in E(R)$ if and only if the lune specified by e contains (strictly within its boundary) fewer than two vertices of X .*

Theorem 2.7 (see [7]) *Let R be the 2-RNG on a given set X , with $|X| = n$. Then*

1. R is 2-connected.
2. R can be constructed in time $O(n^2)$.
3. The number of edges of R is $O(n)$.
4. There exists a $(2, 0)$ -MBSN, say N_0 , on X which is a subgraph of R . If R is given, N_0 can be constructed in a time of $O(n \log n)$.

Let G be a graph embedded in \mathbb{R}^2 and consider the following four variables: k' ; $S_{k'} = \{s_1, \dots, s_{k'}\}$, which is a set of k' distinct Steiner points in \mathbb{R}^2 ; $E_S \subset S_{k'}^2$; and $\mathcal{V} = \{V_1, \dots, V_{k'}\}$, which is a set of subsets of X . Let $Q = \langle V(Q), E(Q) \rangle$ where $V(Q) = X \cup S_{k'}$ and $E(Q) = E(G) \cup E_S \cup \{s_i x_j \mid 1 \leq i \leq k', x_j \in V_i\}$. If Q is 2-connected then we call Q a k -block closure of G . If $\ell_{\max}(Q) \leq \ell_{\max}(Q')$ for any k -block closure Q' of G , then Q is an *optimal k -block closure* of G . Note that there may be many distinct optimal k -block closures for G , and a k -block closure exists for any graph G when $k \geq 2$.

A *Steiner edge* is an edge incident to a Steiner point, and for any graph or vertex set U a *Steiner U -edge* is an edge incident to both $S_{k'}$ and U .

Lemma 2.8 ([6]) *For every leaf-block Y of G there exists at least one Steiner $\text{int}(Y)$ -edge in any k -block closure of G .*

Lemma 2.9 ([6]) *For every isolated block W of G there exists at least two distinct Steiner W -edges in any k -block closure of G .*

In this paper the construction of an optimal k -block closure will usually involve *smallest colour-spanning disks* (SCSDs), where the centre of each disk gives the location of the required Steiner point. Given a partition of a set X into $\{V_i\}$ where each V_i is assigned a unique colour, an SCSD is a circle of minimum radius that contains at least one point of each colour. If $|X| = n$ and $|\{V_i\}|$ is constant then an SCSD C can be found in time $O(n \log n)$; see [1, 3]. Clearly C is determined by either two diametrically opposite points, or by three points. These points are referred to (in [3]) as the *determinators* of C .

One way of constructing SCSDs is by way of a *farthest colour Voronoi diagram* (FCVDs). The FCVD is defined in [1] as follows. Let $\mathcal{C} = \{P_1, \dots, P_q\}$ be a collection of q sets of n

coloured points. If $p \in P_i$, i.e., p is a point of colour i , we put all points of the plane in the *region* of p for which i is the farthest colour, and p the nearest i -coloured point. In other words, x belongs to the region of p if and only if the closed circle centred at x that passes through p contains at least one point of each colour, but no point of colour i is contained in its interior. The FCVD for \mathcal{C} , written $\text{FCVD}(\mathcal{C})$, is the decomposition of the plane into these regions; in other words the edges and vertices of the FCVD are the intersections of boundaries of regions.

Theorem 2.10 ([1]) *For constant q $\text{FCVD}(\mathcal{C})$ can be computed in $O(n^2)$ time, and its structural complexity is $O(n)$.*

Since the centre of an SCSD is either a vertex or the midpoint of an edge of the FCVD we have the following result.

Corollary 2.11 ([1]) *Given the $\text{FCVD}(\mathcal{C})$, an SCSD on \mathcal{C} can be found in $O(n)$ time.*

To conclude this section we state a proposition and two corollaries that allow us to assume a reduced structural complexity for N_k ; namely that N_k has $O(n)$ edges, that the degrees of all vertices of N_k are bounded by a constant, and that the BCF of the subgraph of N_k induced by its non-Steiner edges has a bounded number of leaf-blocks and isolated blocks.

Proposition 2.12 ([6]) *There exists a $(2, k)$ -MBSN N_k on X such that N_k is a subgraph of the 2-RNG on $V(N_k)$ and the degree of v is at most 5 for every $v \in V(N_k)$.*

In the rest of this paper we assume that N_k is a $(2, k)$ -MBSN on X , with Steiner point set $S_{k'}$, satisfying Proposition 2.12. We also assume that there are no chord-paths (i.e., paths with both end-vertices on the same cycle, but sharing no edges with the cycle) in N_k with interiors consisting of degree-two Steiner points only; note that such paths can simply be deleted from N_k without affecting Proposition 2.12.

An *external Steiner edge* is a Steiner edge with one end-point not in $S_{k'}$. Let m_0 be the number of external Steiner edges of N_k . For any G we denote the edge-subgraph of G containing all edges of G of length at most r by $G(r)$. Let R be the 2-RNG on X and let $\overline{N_k} := N_k - S_{k'}$. Clearly $\overline{N_k}$ is a subgraph of $R(\ell_{\max}(\overline{N_k}))$.

Corollary 2.13 ([6]) $b(R(\ell_{\max}(\overline{N_k}))) \leq b(\overline{N_k}) \leq m_0 \leq 5k$

Corollary 2.14 ([6]) *Let $G = R(\ell_{\max}(\overline{N_k}))$ and let G^+ be any optimal k -block closure of G . Then G^+ is a $(2, k)$ -MBSN on X .*

3 The fixed topology on subsets problem

Bae et al. provide new machinery for solving the *bottleneck k -Steiner tree problem*, in other words the $(1, k)$ -MBSN problem, in [2, 3]. Their algorithm depends (as does ours) on being able to find an optimal solution for a given topology, i.e., where the Steiner points are variable but the topology and the terminals are fixed. In fact, they essentially show that it is not even necessary for the terminals in the given topology to be specified in order to find an optimal solution, as long as it is known which subtrees of the minimum spanning tree on X contain the endpoints of the external Steiner edges. They refer to this more general problem as the *fixed topology on subtrees problem*. We show that only a minor reformulation of the problem is needed in order to make this machinery suitable as a component for $(2, k)$ -MBSN construction. We begin by restating the problem in the form originally given by Bae et al.

An *abstract topology* \mathcal{T}_0 is a tree topology on $a' + k + 1$ vertices $s_1, \dots, s_k, t^1, \dots, t^{a'+1}$, where a' is a positive integer, t^i represents a subtree T_i of the minimum spanning tree (MST) on X , and s_j represents a Steiner vertex (i.e., a non-embedded vertex that represents a Steiner point). It is assumed that each s_j is of degree at most five (and no less than 2) and that \mathcal{T}_0 is a full Steiner tree (in other words its terminals are all of degree 1); as Bae et al. pointed out, if t^i is an internal vertex of \mathcal{T}_0 then we can split \mathcal{T}_0 at t^i into two abstract topologies and consider each independently.

Problem 3.1 (The fixed topology on subtrees problem) *Suppose we are given a set X of n terminals in the plane, two positive integers $k \geq 2$ and $a' \leq 4k$, and an abstract topology \mathcal{T}_0 on $s_1, \dots, s_k, t^1, \dots, t^{a'+1}$, where $\{T_i\}$ is the set of subtrees that results by removing a' longest edges from the MST on X . Find an optimal placement of the k Steiner vertices to obtain a minimum bottleneck Steiner tree T^* with the same topology as \mathcal{T}_0 , such that if s_j is adjacent to t^i in \mathcal{T}_0 then s_j is adjacent to a closest vertex of T_i in T^* .*

The method of solution that Bae et al. employ for the problem runs in $O(n^k)$ time and essentially consists of four steps:

1. Enumerate all possible combinations of determinators for the abstract topology \mathcal{T}_0 . This step is achieved by constructing $\text{FCVD}(\mathcal{C}_i)$ for each Steiner point s_i , where $\mathcal{C}_i = \{V(T_j) \mid t_j \text{ adjacent to } s_i \text{ in } \mathcal{T}_0\}$. Using FCVD's in this step prunes the total number of combinations to $O(n^k)$ from $O(n^{3k})$.
2. Each set of determinators from the previous step defines a so called *concrete topology* \mathcal{T} which is a sub-topology of \mathcal{T}_0 , but where the terminals of \mathcal{T} are exactly the members of X that are determinators for \mathcal{T}_0 . As an example consider the abstract topology depicted in Fig. 2, where the t^i are represented by larger white-filled circles and the s_j by smaller black-filled circles. The combination of determinators for this example is such that edges $s_1 t^1$ and $s_1 t^2$ are deleted in order to construct the corresponding concrete topology \mathcal{T} . All solid edges are edges incident to determinators in \mathcal{T} ; the

edge s_2s_3 , although not incident to a determinator, is included in \mathcal{T} since it is an internal edge.

Next, arrows are placed on the edges of \mathcal{T} resulting in a directed tree such that, for any pair of adjacent Steiner points u and v , u is directed to v if and only if v is a determinator of u . A partial ordering on sets of Steiner points is thereby induced by \mathcal{T} which has the property that, for the given set of determinators, $e_{\max}(T^*)$ is always incident to some maximal element of the ordering. The maximal sets of Steiner points are referred to as *primary clusters* in [2]. In Fig. 2 the primary clusters are $\{s_1, s_2\}$ and $\{s_4, s_5, s_6\}$.

3. For each primary cluster $S^i = \{s_j\}$ let \mathcal{T}_i be the subtree of \mathcal{T} induced by S^i and the determinators of the vertices in S^i . A function h_e , which has as input the variable locations of the Steiner vertices of \mathcal{T}_i and outputs the length of $e \in E(\mathcal{T}_i)$ given these locations, is constructed and has the following property. Define a *critical value* of the bottleneck objective function $\ell_{\max}(\mathcal{T}_i)$ as a value \mathbf{q} of the domain of h_e for which $h_{e_1}(\mathbf{q}) = h_{e_2}(\mathbf{q})$ for every pair of edges e_1, e_2 of \mathcal{T}_i . Then every critical value is a vertex of the lower or upper envelope of the set of surfaces $\{h_e : e \in E(\mathcal{T}_i)\}$ (each h_e is in fact a quadratic surface embedded in $2m + 1$ dimensional space, where m is the number of Steiner vertices of \mathcal{T}_i). The set of all critical values across all primary clusters is found in a time bounded above by a function of k only.
4. Finally, a binary search on the collected critical values, using a modified version of the decision algorithm in [13] and incorporating FCVDs, is performed in order to find the optimal solution.




Figure 2: Primary clusters $\{s_1, s_2\}$ and $\{s_4, s_5, s_6\}$

Note that Bae et al.'s solution does not depend on the fact that the T_i are trees, but only that \mathcal{T}_0 is a tree (in order for the direction of \mathcal{T} to induce a partial ordering in Step (2)). Therefore to incorporate their method into our algorithm we will begin by showing that an analogous graph \mathcal{N} , which is the topology of the subgraph of N_k induced by all Steiner edges, contains no cycles consisting of Steiner points only.

We state the next two results for optimal k -block closures, but it is clear that the proofs also hold for N_k . Theorem 3.2 has an analog in the classical *minimum Steiner 2-connected network* problem (see [9]).

Theorem 3.2 *For any G there exists an optimal k -block closure G^+ such that there is no chord-path in G^+ consisting of Steiner points only.*

Proof. Let G_0 be any optimal k -block closure of G . We assume that there are no chord-paths in G_0 with interior vertices that are all degree-two Steiner points nor any non-critical Steiner edges, since any such chord-paths or edges may simply be removed without affecting connectivity. Of all chord-paths of G_0 consisting entirely of Steiner points let P be one containing the least number of edges, and let C be a cycle of which P is a chord-path. Let the end-vertices of P be v, w , let s be a Steiner point of degree at least three in the interior of P , and let u be a neighbour of s not on P . Finally, let P_v, P_w be the subpaths of P that partition the edge-set of P at s . By the 2-connectivity of G_0 there exists a path P' that connects u and v but does not contain s . Regardless of the location of the first intersection of P' with $C \cup P$ a new cycle is produced with a chord-path formed by a subpath of P_v or P_w . This contradicts the fact that P is a chord-path with the least number of edges. ■

Corollary 3.3 *For any G there exists an optimal k -block closure G^+ such that there is no cycle in G^+ consisting of Steiner points only.*

Proof. Any cycle of a 2-connected graph which is not itself a cycle contains a subpath that is a chord of some cycle. ■

We are now ready to update the definition of an abstract topology and to reformulate Problem 3.1. The graph \mathcal{N} is called an *abstract topology* if and only if \mathcal{N} is full Steiner tree topology on $a'' + k'$ vertices $s_1, \dots, s_{k'}, X_1, \dots, X_{a''}$, where $k' \geq 2$, $a'' \leq 5k'$, each X_i is a labelled subset of X , and the degree of any Steiner point is at most 5 and no less than 2. If X_0, X_1 are two distinct terminals of \mathcal{N} then we allow for the possibility that $X_0 \cap X_1 \neq \emptyset$. We observe that this final condition does not affect the feasibility of Bae et al.'s solution since if any two Steiner points share a determinator $x \in X$ then, before constructing the FCVD's in Step (1) above, we simply split x into two distinct terminals occupying the same location. Therefore the solution to the fixed topology on subtrees problem can be utilised with minor modifications (and without a change in complexity) for the following problem.

Problem 3.4 (Fixed topology on subsets problem) *Suppose we are given a set X of n terminals in the plane, a positive integer $k' \geq 2$, and an abstract topology \mathcal{N} . Find a placement of the k' Steiner vertices to obtain a minimum bottleneck Steiner tree N^* that has the same topology as \mathcal{N} , such that if s_j is adjacent (in \mathcal{N}) to X_i then s_j is adjacent (in N^*) to a closest vertex of X_i .*

The role of \mathcal{T}_0 in the construction of minimum $(1, k)$ -MBSNs is clear, since in fact $T^* \cup \bigcup \{T_i\}$ is the required solution when \mathcal{T}_0 is chosen correctly; the graph \mathcal{N} plays a similar role

in our algorithm. We close this section with a simple monotonicity-on-subsets property of solutions to the above problem.

Lemma 3.5 (Monotonicity) *Let $X'_1, \dots, X'_{a''}$ be labelled subsets of X such that $X_i \subseteq X'_i$ for every i , and let N_0^* be a solution to Problem 3.4 when the terminals of \mathcal{N} are the X'_i . Then $\ell_{\max}(N_0^*) \leq \ell_{\max}(N^*)$.*

Proof. Note that N^* is also a bottleneck Steiner tree for \mathcal{N} when its terminals are the X'_i . ■

4 Constructing \mathcal{N} efficiently

By Corollary 2.14 one can construct a $(2, k)$ -MBSN on X by finding an optimal k -block closure of the subgraph $R(d)$ of the 2-RNG R on X , where $d = \ell_{\max}(\overline{N}_k)$. In turn, one can construct an optimal k -block closure for $G = R(d)$ by solving the fixed topology on subsets problem for every abstract topology \mathcal{N} on subsets of X , and choosing a cheapest solution for which $G \cup N^*$ (with N^* as defined in Problem 3.4) is 2-connected. A key part of this latter problem, which we deal with in this section, is to prune the number of potentially optimal combinations of subsets of X which are to serve as the terminals of \mathcal{N} (i.e., to achieve the bound $a'' \leq 5k'$ in Problem 3.4). This is clearly necessary because the total number of subsets of X is 2^n .




Figure 3: A 3-block closure M^3 of M

Consider the graph M^3 in Fig. 3 which is a 3-block closure of the graph M from Fig. 1. In this representation of M^3 we denote Steiner points by black-filled circles and Steiner edges of M^3 by single lines. An edge between a Steiner point s_j and a vertex Y_i of $F_{\mathcal{Y}(M)}$ means that s is adjacent to some vertex of Y_i in M^3 . Now suppose that we modify M^3 by replacing e_1 by an edge e'_1 incident to s_1 and some Y_i with $i \leq 5$, and replacing e_2 by an edge e'_2 incident to s_2 and Y_j for any j . Then the resultant graph will still be 2-connected

whenever $j \geq i$, except in some specific instances when $j = i$ and e'_1 or e'_2 are incident to a cut-vertex of M . We refer to (e_1, e_2) as a *linked pair* of the path $H = Y_1, \dots, Y_6$. In the next subsection we make this notion more precise, and then expound on its utility in the remaining subsections.

4.1 Linked pairs

Let \mathcal{H} be a minimum-cardinality partition of the BCF $F_{\mathcal{Y}(G)}$ such that each member of \mathcal{H} is a path in $F_{\mathcal{Y}(G)}$ with interior vertices of degree-two only. We may assume for simplicity that each member $H \in \mathcal{H}$ is of the form Y_1, \dots, Y_{q_H} for some $q_H \geq 1$, where each Y_i is a block; see Fig. 4. Clearly $q_H = 1$ only if H is an isolated vertex of $F_{\mathcal{Y}(G)}$. Suppose that $q_H \neq 1$. For each $i \in \{2, \dots, q_H\}$ let τ_i be the cut-vertex of G common to Y_i and Y_{i-1} . Let G_0 be the component of $G - \tau_2$ containing no vertices of Y_2 , and let $W_0 = G_0 - \text{int}(Y_1)$. Let $\tau_1 = W_0 \cap Y_1$ (note that τ_1 may contain more than one vertex) and let $\text{int}(W_0) = W_0 - \tau_1$. For every $i \in \{1, \dots, q_H\}$ let $W_i = Y_i - \tau_i$. Finally, let G_1 be the component of $G - \tau_{q_H}$ containing no vertices of Y_1 , let $W_{q_H+1} = G_1 - Y_{q_H}$, and let $\tau_{q_H+1} = Y_{q_H} - \text{int}(Y_{q_H}) - \tau_{q_H}$. Note that W_0 or W_{q_H+1} may be empty. In Fig. 4 we provide an example of some of these definitions with respect to the graph $G = M$. In this example $q_H = 6$, W_1 is a vertex, $\tau_1 = \{\tau_1^1, \tau_1^2\}$, and W_{q_H+1} is empty since Y_6 is a leaf-block.




Figure 4: A partition \mathcal{H} of M , with member H ; the subgraph \hat{H} of M highlighted; and the subgraphs W_0 and W_{q_H} highlighted

Now consider any k -block closure G^k of G . We begin a recursive definition by setting $t = 0$, $\text{rg}(t) = 0$, $H_t = H$, and $W_0^t = W_0$. In general we have the following definition: let P_{t+1} be a path in G^k with no internal vertices in \hat{H} connecting $\text{int}(W_0^t)$ to $W_{\text{rg}(t+1)}$ where $\text{rg}(t+1) \in \{\text{rg}(t) + 1, \dots, q_H + 1\}$ is chosen to be as large as possible. Let $Y^{t+1} = Y_1 \cup \dots \cup Y_{\tilde{t}}$ where $\tilde{t} = \min\{q_H, \text{rg}(t+1)\}$, let $W_0^{t+1} = Y^{t+1} \cup P_1 \cup \dots \cup P_{t+1}$, and let H_{t+1} be the BCF of the graph $\hat{H}_{t+1} = \hat{H} \cup P_1 \cup \dots \cup P_{t+1}$; see Fig. 5.

If W_0 is not empty then P_1 exists lest a vertex in τ_1 be a cut-vertex of G^k . However, if Y_1 is a leaf-block of G then W_0 is empty and we set $\text{rg}(1) = 1$ and $P_1 = \emptyset$. For general

$t > 0$ the path P_{t+1} exists in G^k unless $\text{rg}(t) = q_H + 1$. Let b be the value of t for which $\text{rg}(t+1) = q_H + 1$. Note that for any $t < b - 1$ the end-edge of P_{t+1} incident to $W_{\text{rg}(t+1)}$ is a Steiner edge.

Observation 4.1 W_0^{t+1} is 2-connected. Therefore H_{t+1} is a path and W_0^{t+1} is an end-block of \hat{H}_{t+1} . The complete list of blocks of \hat{H}_{t+1} as they appear in the path are $W_0^{t+1}, Y_{\text{rg}(t+1)+1}, \dots, Y_{q_H}$ where $W_0^{t+1} \cap Y_{\text{rg}(t+1)+1} = \tau_{\text{rg}(t+1)+1}$ (therefore $\text{int}(W_0^{t+1}) = W_0^{t+1} - \tau_{\text{rg}(t+1)+1}$). In particular, \hat{H}_{b+1} is 2-connected.

By the maximality of $\text{rg}(t+1)$ it follows that P_{t+1} is internally disjoint from P_t for all $t > 0$. Therefore, since P_{t+1} is incident to the interior of W_0^t , P_{t+1} has an end-vertex in some $W_{\text{lf}(t)}$ where $\text{rg}(t-1) \leq \text{lf}(t) < \text{rg}(t)$, or in $\text{int}(Y_{\text{rg}(t)})$. For $0 < t \leq b$ let e_2^t be the Steiner edge contained in P_t and incident to $W_{\text{rg}(t)}$ and let e_1^t be the Steiner edge contained in P_{t+1} and incident to $W_{\text{lf}(t)}$. Note that $\text{lf}(t) = \text{rg}(t-1)$ only if e_1^t is incident to $\tau_{\text{rg}(t-1)}$. We refer to (e_1^t, e_2^t) as a *linked pair* of H (see Fig. 5). Each e_1^t is referred to as the *left* member of the linked pair, and e_2^t as the *right* member. Let $E_H = \{(e_1^t, e_2^t)\}$ be the set of all linked pairs of H and observe that $|E_H| = b$.




Figure 5: Linked pair (e_1^t, e_2^t) constructed at step $t+1$

4.2 Viable neighbour-substitutions

The main result of this subsection is Proposition 4.4, which essentially shows that, for any $H \in \mathcal{H}$, the connectivity of G^k only depends on the left/right ordering of the edges in each linked pair incident to \hat{H} (as is the case in the example of Figure 3). For any vertex x of \hat{H} let $\text{in}(x) = j$ (the *index* of x) if and only if W_j contains x for some $j \in \{1, \dots, q_H\}$; we also similarly define $\text{in}(e) = j$ when an end-vertex of e is contained in W_j . The *neighbour*

substitution operation takes as input (e, x) , where e is a Steiner \widehat{H} -edge and x is a vertex of \widehat{H} , and replaces the current non-Steiner end-vertex of e by x .

Let $\mu = \mu(H) = (e_1^1, e_2^1, \dots, e_1^b, e_2^b, f_1, \dots, f_{p(H)})$ where the e_i^j are the members of E_H , and $E_H' = \{f_i : 1 \leq i \leq p(H)\}$ is the set of all Steiner \widehat{H} -edges that do not occur in any member of E_H . Let $\chi = \chi(H) = (\chi^1, \dots, \chi^{2b+p(H)})$ be a sequence of (not necessarily distinct) vertices of \widehat{H} . Let $G^k(\chi)$ be the graph that results from G^k by simultaneously performing all neighbour substitutions contained in $\{(\mu^i, \chi^i)\}$, where μ^i is the i -th term of μ . Let $P_1(\chi), \dots, P_{b+1}(\chi)$ be the paths of $G^k(\chi)$ that result from the paths P_1, \dots, P_{b+1} of G^k after the substitution; in other words each $P_i(\chi)$ is identical to P_i except possibly at the end-vertices. We say that χ is *viable* if and only if the following conditions hold:

- (1) $\text{in}(\chi^2) \geq 1$.
- (2) $\text{in}(\chi^{i-1}) \leq \text{in}(\chi^i)$ and $\chi^{i-1} \neq \tau_{\text{in}(\chi^i)+1}$ for any even $i \leq 2b$.
- (3) χ^i is not a cut-vertex of $G^k(\chi)$ for any i such that $1 \leq i \leq 2b + p(H)$.

Clearly $\chi = \chi_0$ is viable if χ_0 is the sequence of end-vertices of μ in G^k (i.e., $G^k(\chi_0) = G^k$). We now prove a number of results necessary for Proposition 4.4. For simplicity the edge symbols occurring in μ will also be used to denote the corresponding edges in $G^k(\chi)$ for any χ (but the context will always be clear). Also, when e_1^t, e_2^t are considered as edges of $G^k(\chi)$ we use the symbols $\text{rg}(t, \chi)$ and $\text{lf}(t, \chi)$ for $\text{in}(e_2^t)$ and $\text{in}(e_1^t)$ respectively; therefore $\text{rg}(t, \chi_0) = \text{rg}(t)$ and $\text{lf}(t, \chi_0) = \text{lf}(t)$.

Let χ be a sequence satisfying Conditions (1) and (2) of viability, and let $t \in \{1, \dots, b+1\}$. If $t < b+1$ let $U^t = U^t(\chi) = \bigcup_{1 \leq j \leq \widetilde{\text{rg}}} Y_j \cup \bigcup_{j \in I(1,t)} P_j(\chi)$, where $\widetilde{\text{rg}} = \widetilde{\text{rg}}(t, \chi) = \max\{\text{rg}(t', \chi) : 1 \leq t' \leq t\}$ and for any integers $q, q' \in \{1, \dots, b+1\}$, the index-set $I(q, q') \subseteq \{q, \dots, q'\}$ contains all indices j such that the end-vertices of P_j are distinct. Similarly, if $t > 1$ let $\overline{U}^t = \overline{U}^t(\chi) = \bigcup_{\widetilde{\text{lf}} \leq j \leq q_H} Y_j \cup \bigcup_{j \in I(t, b+1)} P_j(\chi)$, where $\widetilde{\text{lf}} = \widetilde{\text{lf}}(t, \chi) = \min\{\text{lf}(t' - 1, \chi) : t \leq t' \leq b+1\}$. Finally, let $U^{b+1} = U^{b+1}(\chi) = \widehat{H} \cup \bigcup_{j \in I(1, b+1)} P_j(\chi)$ and let $\overline{U}^1 = \overline{U}^1(\chi) = U^{b+1}$.

Lemma 4.2 *If χ satisfies Conditions (1) and (2) of viability then U^t and \overline{U}^t , with $1 \leq t \leq b+1$, are 2-connected.*

Proof. We use induction on t to prove the case for U^t ; the other case is similar. Clearly the lemma holds when $b = 0$; so assume that $b > 0$ (i.e., there is at least one linked pair). If $t = 1$ then $U^1 = Y_1 \cup \dots \cup Y_{\text{rg}(1, \chi)} \cup P_1(\chi)$ since e_2^1 is incident to $W_{\text{rg}(1, \chi)}$ (i.e., $\text{rg}(1, \chi) > 0$ and e_2^1 is not incident to $\tau_{\text{rg}(1, \chi)}$). Therefore U^1 is 2-connected by Lemma 2.4.

Next suppose that U^t is 2-connected for some $1 \leq t < b$. Note first that $\text{lf}(t, \chi) \leq \text{rg}(t, \chi) \leq \widetilde{\text{rg}}$, so that if the end-vertices of $P_{t+1}(\chi)$ coincide then they do so at a vertex of U^t . Therefore

in this case $U^{t+1} = U^t$ is 2-connected, and henceforth we assume that the end-vertices of $P_{t+1}(\chi)$ do not coincide. If $\text{rg}(t+1, \chi) \leq \tilde{\text{rg}}$ then by Lemma 2.4 (with $e = e_2^{t+1}$, $Y_1^A = U^t$ and Y_q^A the edge of $P_{t+1}(\chi)$ adjacent to e_2^{t+1}) we conclude that U^{t+1} is 2-connected. Similarly, when $\text{rg}(t+1, \chi) > \tilde{\text{rg}}$ then U^{t+1} is 2-connected since, by Condition (2) of viability, e_1^t is not incident to $\tau_{\tilde{\text{rg}}+1}$. The reasoning when $t = b$ is similar, and therefore the lemma follows. ■

Lemma 4.3 *Suppose that χ is viable and that $G^k(\chi)$ is 2-connected. Now let χ_1 be a viable sequence that differs from χ in exactly one position. Then $G^k(\chi_1)$ is 2-connected.*

Proof. Suppose first that χ and χ_1 differ in position $j = 2b + i$ for some $i > 0$ and suppose that $G^k(\chi_1)$ has connectivity 1. Let $G' = G^k(\chi) - e'_i$ and note that $G' = G^k(\chi_1) - e'_i$. Since $G^k(\chi)$ is 2-connected the BCF of G' is a path Y'_1, \dots, Y'_z , for some $z > 1$; and therefore since U^{b+1} is 2-connected it is contained in either Y'_1 or Y'_z , say (without loss of generality) Y'_1 . The non-Steiner end-vertex of e'_i in $G^k(\chi_1)$ (namely χ_1^j) is contained in \hat{H} , which is contained in U^{b+1} . Therefore by Lemma 2.4 we must have $\chi_1^j = \tau(Y'_1)$. But then χ_1^j is a cut-vertex of $G^k(\chi_1)$, contradicting the fact that χ_1 is viable. Therefore $G^k(\chi_1)$ is 2-connected.

Next suppose that χ and χ_1 differ in position $j = 2t - 1$ for some $1 \leq t \leq b$ (the remaining case is similar). Let $G'' = G^k(\chi) - e_1^t$ and, as before, assume that $G^k(\chi_1)$ has connectivity 1. The BCF of G'' is a path, say Y''_1, \dots, Y''_z for some $z > 1$, with $U^t(\chi) = U^t(\chi_1)$ a subset of Y''_1 . Since χ_1 is viable $\text{lf}(t, \chi_1) = \text{in}(\chi_1^j) \leq \text{in}(\chi_1^{j+1}) = \text{rg}(t, \chi_1) \leq \tilde{\text{rg}}(t, \chi_1)$ and therefore $\chi_1^j \in Y''_1$. As in the previous case χ_1^j is a cut-vertex of $G^k(\chi_1)$, contradicting the fact that χ_1 is viable. Therefore $G^k(\chi_1)$ is 2-connected. ■

Proposition 4.4 *If χ is viable then $G^k(\chi)$ is 2-connected.*

Proof. We use the previous lemma inductively with χ_0 as the base. ■

4.3 Verifying viability

Our main algorithm contains a search procedure that iterates through viable sequences, and therefore we need to show that a sequence χ can be verified as being viable in “reasonable” time. Under the assumption that preprocessing has been performed on G , specifically that the BCF of G has been specified, it should be clear that Conditions (1) and (2) of viability are verifiable in $O(k)$ steps. We now show that Condition (3) can also be verified in a total number of steps bounded above by a function of k only.

Let χ_1 be any neighbour-substitution sequence that satisfies Conditions (1) and (2) for viability, but not Condition (3). Since the only edges of $G^k(\chi_1)$ incident to $U^{b+1}(\chi_1)$ are Steiner edges or edges incident to τ_1 or τ_{q_H+1} , and $U^{b+1}(\chi_1)$ is 2-connected, any cut-vertex of $G^k(\chi_1)$ at some χ_1^j separates $U^{b+1}(\chi_1)$ from some subgraph of $G^k(\chi_1)$. This means that any minimal edge-cut of $G^k(\chi_1)$ may only contain Steiner edges, or edges incident to

τ_1 or τ_{q_H} and not contained in $E(\widehat{H})$. We now see that for every Steiner edge $e \in E(G^k)$ incident to \widehat{H} there exists a unique set $\beta(e)$ such that either $\beta(e) = \emptyset$ or $e \in \beta(e)$ and, for any sequence χ_1 satisfying Conditions (1) and (2) but such that the endpoints of all edges in $\beta(e)$ coincide, Condition (3) is violated and $\beta(e)$ is a minimal edge-cut of $G^k(\chi_1)$. As an example suppose that G^k is the graph depicted in Fig. 6. Let τ' be the cut-vertex of G common to blocks Z and Y_1 , and let E_0 be the set of all edges in Z that are incident to τ' . Then $\beta(f_1) = \{f_1\} \cup E_0$ and $\beta(e_1^1) = \beta(e_2^1) = \emptyset$. We illustrate another example in Fig. 7 where this time $H = Y_1, \dots, Y_5$. Then $\beta(e_2^2) = \beta(e_1^1) = \{e_1^1, e_2^2\}$ and $\beta(e_1^2) = \beta(e_2^1) = \emptyset$.

From the above discussion the $\beta(e)$ are independent of the choice of χ_1 , and can be constructed once the topology of \mathcal{N} and the BCF of G are known. We therefore refer to $\{\beta(e)\}$ as the set of *potential cuts*. Furthermore, in the main algorithm we only consider G and \mathcal{N} with $b(G) \leq 5k$ and $|E(\mathcal{N})| \in O(k)$. Therefore $|\mathcal{H}| \in O(k)$ and all $\beta(e)$ (for all $H \in \mathcal{H}$) can be constructed in the required time; hence Condition (3) can also be verified within this time for any neighbour-substitution sequence.




Figure 6: $\beta(f_1) = \{f_1\} \cup E_0$, $\beta(e_1^1) = \beta(e_2^1) = \emptyset$




Figure 7: $\beta(e_2^2) = \beta(e_1^1) = \{e_1^1, e_2^2\}$, $\beta(e_1^2) = \beta(e_2^1) = \emptyset$

4.4 Constructing an abstract topology by means of linked pairs

In this subsection we show how the choice of linked pairs and the topology induced by the Steiner edges of G^k may be used to construct an abstract topology. Let \mathcal{N} have the afore-mentioned topology. The terminals of \mathcal{N} are specified by partitioning each \widehat{H} into either one or two colour sets for each Steiner edge incident to \widehat{H} . In particular, for every linked pair incident to \widehat{H} we partition \widehat{H} at a cut-vertex separating the end-vertices of the pair; and for every non-linked edge we essentially colour the whole of \widehat{H} by a single colour. The advantage of using this method becomes apparent in Proposition 4.6, where it is shown that if G , \mathcal{N} , and the colour partitions are chosen correctly, then a solution to the fixed topology on subsets problem for \mathcal{N} yields a $(2, k)$ -MBSN on X .

More formally, let $E_S = \{e_j\}_{j \in I_S}$, for some index set I_S , be the set of external Steiner edges of G^k . Let the *marker* of any linked edge e_j^i of G^k initially be defined as $\text{mk}(e_j^i) := \text{in}(e_1^i)$, and let $\mathcal{M} = \{\text{mk}(e) : e \text{ is a linked edge of } G^k\}$. For any $e \in E_S$ incident to some $\widehat{H}(e)$, where $H(e) \in \mathcal{H}$ and $H(e) = Y_1, \dots, Y_{q_{H(e)}}$, we define the *colour-set* of e with respect to \mathcal{M} as

$$J(e, \mathcal{M}) = \begin{cases} Y_1 \cup \dots \cup Y_{\text{mk}(e) - \tau_{\text{mk}(e)+1}} & \text{if } e \text{ is a left linked edge} \\ W_{\text{mk}(e)} \cup \dots \cup W_{q_{H(e)}} & \text{if } e \text{ is a right linked edge} \\ Y_1 \cup \dots \cup Y_{q_{H(e)}} & \text{for all other } e \end{cases}$$

For any G' let $S(G')$ be the subgraph of G' induced by the Steiner edges. Let \mathcal{G} be the set of all k -block closures of G such that for any $G' \in \mathcal{G}$ the graph $S(G')$ is isomorphic to \mathcal{N} , and for any $j \in E_S$ edge e_j is incident to a vertex of $J(e_j, \mathcal{M})$ in G' . Observe that $G^k \in \mathcal{G}$. We begin by letting the terminals of \mathcal{N} be the members of $\{J(e_j, \mathcal{M})\}_{j \in I_S}$; specifically, for every $j \in I_S$ the terminal incident to edge e_j is $J(e_j, \mathcal{M})$. If e_j and e_i are distinct members of E_S then $J(e_j, \mathcal{M})$ and $J(e_i, \mathcal{M})$ are considered as distinct terminals of \mathcal{N} , even if they contain exactly the same vertices (that is, we think of the terminals as labelled subsets of X). Therefore, from the discussion in Section 3, we may assume that \mathcal{N} is a forest with no internal terminals.

Let N^* be a solution to the fixed topology on subsets problem for \mathcal{N} . Note that $\ell_{\max}(G \cup N^*) \leq \min\{\ell_{\max}(G') : G' \in \mathcal{G}\}$ by the monotonicity of N^* , but that $G \cup N^*$ is not necessarily 2-connected. To see this consider the neighbour-sequence $\chi(H)$, for some H , that corresponds to the respective end-vertices of the Steiner edges incident to \widehat{H} in $G \cup N^*$. Then by the definition of $J(e, \mathcal{M})$ we know that $\chi(H)$ satisfies Conditions (1),(2) of viability but not necessarily Condition (3), so that $G \cup N^*$ may contain a cut-vertex. However, we show that there exists a set $\Gamma = \{\gamma(e_j) : j \in I_S\}$, where each $\gamma(e_j)$ is a (possibly empty) subset of $J(e_j, \mathcal{M})$, such that if the terminal-set of \mathcal{N} is modified to $\{J(e_j, \mathcal{M}) - \gamma(e_j)\}_{j \in I_S}$ then $G \cup N^*$ will be an optimal member of \mathcal{G} . We denote this updated abstract topology by $\mathcal{N}(\Gamma)$, and its optimal embedding by $N^*(\Gamma)$. Recall that we may assume G^k contains no chord-paths consisting of Steiner points only.

Lemma 4.5 *There exists a set $\Gamma = \{\gamma(e_j)\}_{j \in I_S}$, where $\gamma(e_j) \subset J(e_j, \mathcal{M})$, such that $G \cup N^*(\Gamma)$ is an optimal member of \mathcal{G} .*

Proof. As before, let $\{\beta(e)\}$ be the set of potential cuts. Note first that $\beta(e)$ does not contain any two edges of the same linked pair, and that no two edges of $\beta(e)$ are contained in the same component of \mathcal{N} unless the component contains no Steiner points of degree more than 2; for otherwise, since U^{b+1} is 2-connected, there would be a chord-path in G^k consisting of Steiner points only.

Let $\{\mathcal{N}^i\}$ be the set of connected components of \mathcal{N} , and let $\{N_i^*\}$ be the corresponding subgraphs of N^* . Let $\{g(i, j)\}$ be the set of external Steiner edges of \mathcal{N}^i , and let $v(i, j)$ be the non-Steiner end-vertex of $g(i, j)$ in N_i^* . We claim that there exists an optimal member of \mathcal{G} , say G_{opt} , with the following property: G_{opt} contains a pair of edges $g(i_1, j_1), g(i_2, j_2)$, where $i_1 \neq i_2$ and both edges are members of $\beta(g(i_1, j_1))$, such that (1) the end-vertex of $g(i_1, j_1)$ in G_{opt} is $v(i_1, j_1)$, and (2) the end-vertex of $g(i_2, j_2)$ in G_{opt} is not $v(i_1, j_1)$. To prove the claim let G_0 be any optimal member of \mathcal{G} and suppose first that (1) is false. Let $\beta(e)$ be any potential cut containing only Steiner edges and let $g(i_1, j_1)$ be any member of $\beta(e)$. We perform the neighbour substitution $(g(i_1, j), v(i_1, j))$ for every j , and let the resultant graph be G'_0 after solving the fixed topology on subsets problem for the modified \mathcal{N}^{i_1} . Note that G'_0 is 2-connected and $\ell_{\max}(G'_0) = G_0$. Note that (2) cannot be false for any member of \mathcal{G} , or else $v(i_1, j_1)$ would be a cut-vertex. Therefore the claim holds.

Now let $\beta(e)$ be any potential cut containing only Steiner edges and let $g(i_1, j_1), g(i_2, j_2)$ be any two members of $\beta(e)$. Let $\gamma(g(i_1, j_1)) = v(i_1, j_1)$ and let $\gamma(g(i_2, j_2)) = J(g(i_2, j_2), \mathcal{M}) - v(i_1, j_1)$. Note that the previous claim still holds for the subset of \mathcal{G} (if non-empty) restricted to graphs where the colour sets of $g(i_1, j_1), g(i_2, j_2)$ have been modified by $\gamma(g(i_1, j_1))$ and $\gamma(g(i_2, j_2))$ respectively. We therefore select any other potential cut $\beta(e')$, distinct from $\beta(e)$, containing only Steiner edges and choose any two members $g(i_3, j_3), g(i_4, j_4)$ as before. We proceed in this way until γ has been defined for a pair of edges from every such potential cut. For every $\beta(e)$ containing a non-Steiner edge, and every Steiner edge $e \in \beta(e)$, let $\gamma(e)$ be the common end-vertex of all edges in $\beta(e)$. For any remaining Steiner edges e , $\gamma(e) := \emptyset$. The graph $G \cup N^*(\Gamma)$ (if it exists) will clearly belong to \mathcal{G} , and a cheapest solution is found by considering every possible combination and ordering of pairs of edges in potential cuts. ■

If G is given then the set \mathcal{G} , and the set Γ satisfying the previous lemma, depend only on \mathcal{N} and \mathcal{M} . We therefore let $G(\mathcal{M}, \mathcal{N})$ be an optimal member of \mathcal{G} produced by the method in the previous lemma.

Proposition 4.6 *There exists an \mathcal{M} and an abstract topology \mathcal{N} on $\{J(e_j, \mathcal{M}) - \gamma(e_j)\}_{j \in I_S}$ such that $G(\mathcal{M}, \mathcal{N})$ is a $(2, k)$ -MBSN on X .*

Proof. Let G^+ be any optimal k -block closure of $G = R(d)$, where $d = \ell_{\max}(\overline{N}_k)$. Let \mathcal{M} be defined on a set of linked pairs of G^+ and let \mathcal{N} be have the topology of $S(G^+)$. The proposition now follows from the previous lemma. ■

5 The algorithm

In order to utilise Proposition 4.6 in our algorithm we need a procedure for constructing every potentially optimal \mathcal{N} . To specify the terminals of \mathcal{N} for a given \mathcal{M} we need the following: the set \mathcal{H} ; a unique $H(e) \in \mathcal{H}$ for each external Steiner edge e , such that e is required to be incident to $\widehat{H(e)}$ in $G(\mathcal{M}, \mathcal{N})$; and the set of linked pairs E_H for every $H \in \mathcal{H}$. We encapsulate this information as follows: let \mathcal{N}_0 be any abstract topology on \mathcal{H} and let $\mathcal{E} = \{E_H : H \in \mathcal{H}\}$. Then $(\mathcal{N}_0, \mathcal{E})$ is an instance of an *edge-linked abstract topology* on G .

An *optimal embedding* of $(\mathcal{N}_0, \mathcal{E})$ is a cheapest $G(\mathcal{M}, \mathcal{N})$ for all marker-sets \mathcal{M} and abstract topologies \mathcal{N} consistent with \mathcal{E} and \mathcal{N} . We now present a binary search method, acting recursively on the possible locations of the markers, to find an optimal embedding of $(\mathcal{N}_0, \mathcal{E})$.

Let $a = 1$ and for each linked Steiner edge e of \mathcal{E} let $\text{mk}_a(e)$ be a median of $I_{H(e)} = \{1, \dots, q_{H(e)}\}$, and let \mathcal{M}_a be the set of these markers. For any Steiner point s of $G(\mathcal{M}_a, \mathcal{N})$ we place an arbitrary ordering $e^1, \dots, e^{p(s)}$, where $p(s) \leq \deg(s)$, on the edges incident to s which are members of linked pairs and initially set $j(s) = p(s)$. Now let $e_a = e_{\max}(G(\mathcal{M}_a, \mathcal{N}))$. If e_a is not a member of a linked-pair then $G(\mathcal{M}_a, \mathcal{N})$ is an optimal embedding of $(\mathcal{N}_0, \mathcal{E})$. Otherwise suppose that $e_a \in \{e^1, \dots, e^{p(s)}\}$ for some s . For simplicity we assume throughout that $e^{j(s)}$ is a left edge.

Let $a = 2$ and let $\text{mk}_a(e^{j(s)})$ be a median value of $\{\text{mk}_{a-1}(e^{j(s)}), \dots, q_{H(e^{j(s)})}\}$; all other markers remain at their current positions. Reconstruct $G(\mathcal{M}_a, \mathcal{N})$ with respect to the new marker set \mathcal{M}_a and let $e_a = e_{\max}(G(\mathcal{M}_a, \mathcal{N}))$. Now suppose that we are at a general step $a \geq 2$, and suppose that $e_a \in \{e^1, \dots, e^{p(s)}\}$ for some s . Let $j(s)$ be the largest element of $\{1, \dots, p(s)\}$ such that $\text{mk}_a(e^{j(s)}) < q_{H(e^{j(s)})}$ (if no such $j(s)$ exists then we have found an optimal embedding of $(\mathcal{N}_0, \mathcal{E})$). We increment a by 1 and let $\text{mk}_a(e^{j(s)})$ be a median of $\{\text{mk}_{a-1}(e^{j(s)}), \dots, q_{H(e^{j(s)})}\}$. Each $\text{mk}_a(e^i)$, where $i > j(s)$, is now moved back to a median of $\{1, \dots, q_{H(e^i)}\}$ and $G(\mathcal{M}_a, \mathcal{N})$ is constructed as before. We continue until no further improvement in $\ell_{\max}(G(\mathcal{M}_a, \mathcal{N}))$ can be found. This process of finding an optimal embedding for a given $(\mathcal{N}_0, \mathcal{E})$ is referred to as the *binary marker-search* process.

Let G^{mk} be an optimal k -block closure of G produced by the binary marker-search process. Let $r(G)$ be the length of a longest Steiner edge in G^{mk} . The following result is straightforward (a similar result appears in [6]).

Lemma 5.1 *If G_1 is an edge-subgraph of G_2 then $r(G_1) \geq r(G_2)$.*

A *feasible* topology is 2-connected edge-linked abstract topology $(\mathcal{N}_0, \mathcal{E})$ with $k' \leq k$ Steiner points where the degree of any Steiner points is at most 5.

Theorem 5.2 *Algorithm 1 correctly computes a k -MBSN in a time of $O(n^k \log^{5k-1} n)$.*

Proof. Let $d_{\text{opt}} = \ell_{\max}(\overline{N_k})$ and $G_{\text{opt}} = R(d_{\text{opt}})$. Then by Corollary 2.14 $G_{\text{opt}}^{\text{mk}}$ is a $(2, k)$ -MBSN on X . Any $d \in L$ such that $b(G_d) \leq 5k$, and G_d^{mk} is a $(2, k)$ -MBSN on X is referred

Algorithm 1 Construct a $(2, k)$ -MBSN

Input: A set X of n vertices embedded in the Euclidean plane and a positive integer k

Output: A $(2, k)$ -MBSN on X

- 1: Construct the 2-RNG R on X
 - 2: Let L be the ordered set of edge-lengths occurring in R , where ties have been broken randomly
 - 3: Let d be a median of L
 - 4: **repeat**
 - 5: Construct the BCF of $G_d = R(d)$
 - 6: **if** $b(G_d) > 5k$ **then**
 - 7: Exit the loop and let d be the median of the next larger interval of L
 - 8: **for all** feasible topologies $(\mathcal{N}_0, \mathcal{E})$ on G_d **do**
 - 9: Run the binary marker-search procedure on $(\mathcal{N}_0, \mathcal{E})$ and let $r(G_d)$ be the length of the longest Steiner edge in the optimal embedding of $(\mathcal{N}_0, \mathcal{E})$.
 - 10: **if** $r(G_d) \leq t$ **then**
 - 11: Let d be the next smaller median
 - 12: **else**
 - 13: Let d be the next larger median
 - 14: **until** no smaller value of $\max\{r(G_d), d\}$ can be found
 - 15: Output the instance with the minimum $\max\{r(G_d), d\}$
-

to as *valid*. Now let $d \in L$ be some value considered in the binary search. If $b(G_d) > 5k$ then, by Lemma 2.5, there exists a valid d' such that $d' > d$. If $r(G_d) \leq d$ then clearly there exists a valid d' such that $d' \leq d$, and if $r(G_d) > d$ then, by Lemma 5.1, there exists a valid d' such that $d' \geq t$. Therefore a valid d' will be located by the binary search by decreasing d if $r(G_d) \leq d$ and G_d is connected, and increasing d otherwise.

To prove the complexity note that Line 1 requires $O(n^2)$ time; Line 5 requires $O(n)$ time; and Line 8 requires constant time for fixed k since $(\mathcal{N}_0, \mathcal{E})$ has structural complexity of $O(k)$. The binary marker-search procedure in Line (8) performs at most $O(\ln^{5k-2} n)$ main iterations, since a binary search is performed on at most $5k - 2$ Steiner edges (note that at least two Steiner edges will remain adjacent to fixed leaf-blocks). For each iteration within the binary marker-search procedure the fixed topology on subsets problem is solved a constant number of times. This gives a total of $O(n^k \ln^{5k-2} n)$ steps for the binary marker-search. Finally, we note that the repeat-loop in Line 4 requires at most $O(\ln n)$ steps since this entails a binary search on d . The result follows. ■

6 Conclusion

This paper introduces the first polynomial-time (for constant $k > 2$) exact algorithm for constructing $(2, k)$ -MBSN's. At the heart of our algorithm is a procedure that constructs colour-sets consisting of blocks of G separated by linked pairs. A binary search on the markers for each such linked pair yields an optimal colouring of G . Using the solution

method of Bae et al. in [2] for the fixed topology on subsets problem, the colouring of G produces a $(2, k)$ -MBSN. Similarly to [6], it is possible to generalise our algorithm to other L_p planes. In future work we plan to extend these methods even further for the construction of (c, k) -MBSNs for general $c > 2$.

References

- [1] M. Abellanas, F. Hurtado, C. Icking, R. Klein, E. Langetepe, L. Ma, B. Palop and V. Sacristan, The farthest color Voronoi diagram and related problems, Technical Report 002, Institut für Informatik I, Rheinische Friedrich-Wilhelms-Universität, Bonn, 2006.
- [2] S.W. Bae, S. Choi, C. Lee and S. Tanigawa, Exact algorithms for the bottleneck Steiner tree problem (Extended Abstract). *The 20th International Symposium on Algorithms and Computation (ISAAC 2009)*, LNCS 5878, Hawaii, USA, December 2009; 24–33.
- [3] S.W. Bae, C. Lee and S. Choi, On exact solutions to the Euclidean bottleneck Steiner tree problem, *Information Processing Letters* 110 (2010) 672–678.
- [4] P. Berman and A.Z. Zelikovsky, On approximation of the power- p and bottleneck Steiner trees. In: Du, D., Smith, J.M., Rubinstein, J.H. (eds.) *Advances in Steiner trees*, pp. 117–135. Kluwer Academic Publishers, Netherlands (2000).
- [5] M. Brazil, C.J. Ras, K. Swanepoel and D.A. Thomas, Generalised k -Steiner tree problems in normed planes, Submitted for publication, arXiv:1111.1464v1 [math.CO].
- [6] M. Brazil, C.J. Ras, D.A. Thomas, The bottleneck 2-connected k -Steiner network problem for $k \leq 2$, Submitted for publication, arXiv:1108.3655v1 [math.CO].
- [7] M.S. Chang, C.Y. Tang and R.C.T. Lee, Solving the Euclidean bottleneck biconnected edge subgraph problem by 2-relative neighborhood graphs, *Discrete Applied Mathematics* 39 (1992) 1–12.
- [8] G.A. Dirac, Minimally 2-connected graphs, *J. Reine Angew. Math.* 228 (1967) 204–216.
- [9] E.L. Luebke and J.S. Provan, On the structure and complexity of the 2-connected Steiner network problem in the plane, *Operations Research Letters* 26 (2000) 111–116.
- [10] G.S. Manku, A linear time algorithm for the Bottleneck Biconnected Spanning Subgraph problem, *Information Processing Letters* 59 (1996) 1–7.
- [11] R.G. Parker and R.L. Rardin, Guaranteed performance heuristics for the bottleneck traveling salesman problem, *Operations Research Letters* 2 (1984) 269–272.
- [12] A.P. Punnen and K.P.K. Nair, A fast and simple algorithm for the bottleneck biconnected spanning subgraph problem, *Information Processing Letters* 50 (1994) 283–286.
- [13] M. Sarrafzadeh and C.K. Wong, Bottleneck Steiner trees in the plane, *IEEE Trans. Comput.*, 41 (1992), pp. 370–374.

- [14] R. Tarjan, Depth first search and linear graph algorithms, *SIAM Journal of Computing* 1 (1972) 146–160.
- [15] E.A. Timofeev, Minimax 2-connected subgraphs and the bottleneck traveling salesman problem, *Cybernetics and Systems Analysis*, Volume 15, Number 4, 516–521.
- [16] L. Wang and D.Z. Du, Approximations for a bottleneck Steiner tree problem, *Algorithmica*, 2002, **32**, pp. 554–561.